



US006393497B1

(12) **United States Patent**
Arnold et al.

(10) **Patent No.: US 6,393,497 B1**
 (45) **Date of Patent: May 21, 2002**

(54) **DOWNLOADABLE SMART PROXIES FOR PERFORMING PROCESSING ASSOCIATED WITH A REMOTE PROCEDURE CALL IN A DISTRIBUTED SYSTEM**

(75) **Inventors:** Kenneth C. R. C. Arnold, Lexington; James H. Waldo, Dracut; Robert Scheffler, Somerville; Ann M. Wollrath, Groton, all of MA (US)

(73) **Assignee:** Sun Microsystems, Inc., Palo Alto, CA (US)

(*) **Notice:** Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) **Appl. No.:** 09/044,930

(22) **Filed:** Mar. 20, 1998

(51) **Int. Cl.⁷** G06F 13/00

(52) **U.S. Cl.** 709/330; 709/315

(58) **Field of Search** 709/300, 303, 709/310, 315, 330

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,491,946 A	1/1985	Kryskow, Jr. et al.
4,713,806 A	12/1987	Oberlander et al.
4,809,160 A	2/1989	Mahon et al.
4,823,122 A	4/1989	Mann et al.
4,939,638 A	7/1990	Stephenson et al.
4,956,773 A	9/1990	Saito et al.
5,088,036 A	2/1992	Ellis et al.
5,109,486 A	4/1992	Seymour
5,187,787 A	2/1993	Skeen et al.
5,218,699 A	6/1993	Brandle et al.
5,257,369 A	10/1993	Skeen et al.
5,293,614 A	3/1994	Ferguson et al.

5,297,283 A	3/1994	Kelly, Jr. et al.	
5,307,490 A	4/1994	Davidson	et al.
5,311,591 A	5/1994	Fischer	
5,339,435 A	8/1994	Lubkin et al.	
5,386,568 A	1/1995	Wold et al.	

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

EP	0 300 516 A2	1/1989
EP	0 351 536 A3	1/1990
EP	0 384 339 A3	8/1990
EP	0 472 874 A1	3/1992

(List continued on next page.)

OTHER PUBLICATIONS

Sun Microsystems, Java Remote Method Invocation Specification, Feb. 10, 1997, pp. 18-23.*

Aldrich et al., "Providing Easier Access to Remote Objects in Client-Server Systems," System Sciences, 1998, Proceedings of the 31st Hawaii Internat'l. Conference, Jan. 6-9, 1998, pp. 366-375.

(List continued on next page.)

Primary Examiner—Kenneth R. Coulter

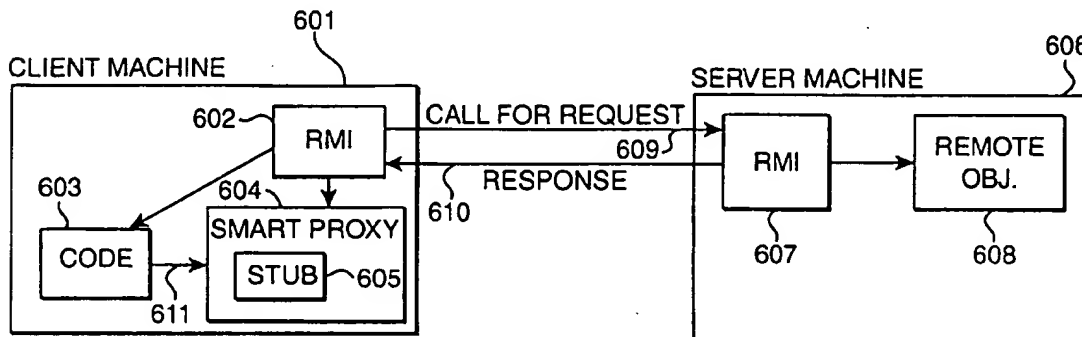
(74) *Attorney, Agent, or Firm*—Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P.

(57) **ABSTRACT**

Use of a smart proxy as a wrapper around a stub in a distributed system. Instead of receiving a stub as a result of a remote procedure call, a caller receives a smart proxy including the stub as an embedded object. The smart proxy performs predefined processing associated with a remote procedure call, the processing possibly occurring before, during, or after a response to the call.

18 Claims, 8 Drawing Sheets

600



U.S. PATENT DOCUMENTS

5,390,328	A	2/1995	Frey et al.	
5,423,042	A	6/1995	Jalili et al.	
5,440,744	A	8/1995	Jacobson et al.	
5,448,740	A	9/1995	Kiri et al.	
5,452,459	A	9/1995	Drury et al.	
5,455,952	A	10/1995	Gjovaag	
5,471,629	A	11/1995	Risch	
5,475,792	A	12/1995	Stanford et al.	
5,475,817	A	12/1995	Waldo et al.	
5,481,721	A	1/1996	Serlet et al.	
5,504,921	A	4/1996	Dev et al.	
5,511,197	A	4/1996	Hill et al.	
5,524,244	A	6/1996	Robinson et al.	
5,548,726	A	8/1996	Pettus	
5,553,282	A	9/1996	Parrish et al.	
5,555,367	A	9/1996	Premarlani et al.	
5,555,427	A *	9/1996	Aoe et al.	709/201
5,557,798	A	9/1996	Skeen et al.	
5,560,003	A	9/1996	Nilsen et al.	
5,561,785	A	10/1996	Blandy et al.	
5,577,231	A	11/1996	Scalzi et al.	
5,594,921	A	1/1997	Pettus	
5,603,031	A	2/1997	White et al.	
5,617,537	A	4/1997	Yamada et al.	
5,628,005	A	5/1997	Hurvig	
5,640,564	A	6/1997	Hamilton et al.	
5,644,768	A	7/1997	Periwal et al.	
5,652,888	A	7/1997	Burgess	
5,655,148	A	8/1997	Richman et al.	
5,659,751	A	8/1997	Heninger	
5,671,225	A	9/1997	Hooper et al.	
5,675,796	A	10/1997	Hodges et al.	
5,680,573	A	10/1997	Rubin et al.	
5,680,617	A	10/1997	Gough et al.	
5,684,955	A	11/1997	Meyer et al.	
5,689,709	A	11/1997	Corbett et al.	
5,706,435	A	1/1998	Barbará et al.	
5,706,502	A	1/1998	Foley et al.	
5,724,588	A	3/1998	Hill et al.	
5,727,145	A	3/1998	Nessett et al.	
5,737,607	A	4/1998	Hamilton et al.	
5,745,678	A	4/1998	Herzberg et al.	
5,745,695	A	4/1998	Gilchrist et al.	
5,745,703	A	4/1998	Cejtin et al.	
5,745,755	A	4/1998	Covey	
5,748,897	A *	5/1998	Katiyar	709/219
5,754,849	A	5/1998	Dyer et al.	
5,757,925	A	5/1998	Faybishenko	
5,761,656	A	6/1998	Ben-Shachar	
5,764,897	A	6/1998	Khalidi	
5,768,532	A	6/1998	Megerian	
5,774,551	A	6/1998	Wu et al.	
5,778,187	A	7/1998	Monteiro et al.	
5,778,228	A	7/1998	Wei	
5,778,368	A	7/1998	Hogan et al.	
5,787,425	A	7/1998	Bigus	
5,787,431	A	7/1998	Shaughnessy	
5,790,548	A	8/1998	Sistanizadeh et al.	
5,802,367	A	9/1998	Held et al.	
5,808,911	A	9/1998	Tucker et al.	
5,809,507	A	9/1998	Cavanaugh, III	
5,812,819	A	9/1998	Rodwin et al.	
5,813,013	A	9/1998	Shakib et al.	
5,815,149	A	9/1998	Mutschler, III et al.	
5,815,709	A	9/1998	Waldo et al.	
5,815,711	A	9/1998	Sakamoto et al.	
5,818,448	A *	10/1998	Katiyar	345/355
5,829,022	A	10/1998	Watanabe et al.	
5,832,219	A	11/1998	Pettus	
5,832,529	A	11/1998	Wollrath et al.	
5,832,593	A	11/1998	Wurst et al.	
5,835,737	A	11/1998	Sand et al.	
5,842,018	A	11/1998	Atkinson et al.	
5,844,553	A	12/1998	Hao et al.	
5,845,129	A	12/1998	Wendorf et al.	
5,860,004	A	1/1999	Fowlow et al.	
5,860,153	A	1/1999	Matena et al.	
5,864,862	A	1/1999	Kriens et al.	
5,864,866	A	1/1999	Henckel et al.	
5,872,928	A	2/1999	Lewis et al.	
5,872,973	A *	2/1999	Mitchell et al.	709/332
5,875,335	A	2/1999	Beard	
5,878,411	A	3/1999	Borroughs et al.	
5,884,024	A	3/1999	Lim et al.	
5,884,079	A	3/1999	Furusawa	
5,887,134	A	3/1999	Ebrahim	
5,889,951	A	3/1999	Lombardi	
5,890,158	A	3/1999	House et al.	
5,892,904	A	4/1999	Atkinson et al.	
5,933,497	A	8/1999	Beetcher et al.	
5,935,249	A	8/1999	Stern et al.	
5,940,827	A	8/1999	Hapner et al.	
5,944,793	A	8/1999	Islam et al.	
5,946,485	A	8/1999	Weeren et al.	
5,946,694	A	8/1999	Copeland et al.	
5,956,509	A	9/1999	Kevner	
5,966,531	A	10/1999	Skeen et al.	
5,969,967	A	10/1999	Aahlad et al.	
5,987,506	A	11/1999	Carter et al.	
5,999,179	A	12/1999	Kekic et al.	
6,003,763	A	12/1999	Gallagher et al.	
6,009,103	A	12/1999	Woundy	
6,016,496	A	1/2000	Roberson	
6,026,414	A *	2/2000	Anglin	707/204
6,031,977	A	2/2000	Pettus	
6,061,699	A	5/2000	DiCecco et al.	
6,061,713	A	5/2000	Bharadhwaj	
6,223,217	B1 *	4/2001	Pettus	709/219
6,260,076	B1 *	7/2001	Savitzky et al.	709/315

FOREIGN PATENT DOCUMENTS

EP	0 474 340	A2	3/1992
EP	0 497 022	A1	8/1992
EP	0 555 997	A2	8/1993
EP	0 565 849	A2	10/1993
EP	0 569 195	A3	11/1993
EP	0 625 750	A2	11/1994
EP	0 635 792	A2	1/1995
EP	0 651 328	A1	5/1995
EP	0 660 231	A2	6/1995
EP	0 697 655	A2	2/1996
EP	0 718 761	A1	6/1996
EP	0 767 432	A2	4/1997
EP	0 778 520	A2	6/1997
EP	0 794 493	A2	9/1997
EP	0 803 810	A2	10/1997
EP	0 803 811	A	10/1997
EP	0 805 393	A2	11/1997
EP	0 810 524	A	12/1997
EP	0 817 020	A	1/1998
EP	0 817 022	A2	1/1998
EP	0 817 025	A	1/1998
EP	0 836 140	A2	4/1998
GB	2 253 079	A	8/1992
GB	2 262 825	A	6/1993
GB	2 305 087	A	3/1997
JP	11-45187		2/1999

WO	WO 92/07335	4/1992
WO	WO 92/09948	6/1992
WO	WO 93/25962 A1	12/1993
WO	WO 94/03855	2/1994
WO	WO 96/03692 A1	2/1996
WO	WO 96/10787	4/1996
WO	WO 96/18947	6/1996
WO	WO 96/24099	8/1996
WO	WO 98/02814	1/1998
WO	WO 98/04971	2/1998

OTHER PUBLICATIONS

- Aldrich et al., "Providing Easier Access to Remote Objects in Distributed Systems," Calif. Institute of Technology, www.cs.caltech.edu/%7Ejedi/paper/jedipaper.html, Nov. 21, 1997.
- Burns et al., "An Analytical Study of Opportunistic Lease Renewal," Distributed Computing Systems, 21st International Conference, pp. 146-153, Apr. 2000.
- Dollimore et al., "The Design of a System for Distributing Shared Objects," The Computer Journal, No. 6, Cambridge, GB, Dec. 1991.
- Fleisch et al., "High Performance Distributed Objects Using Distributed Shared Memory & Remote Method Invocation," System Sciences, 1998, Proceedings of the 31st Hawaii International Conference, Jan. 6-9, 1998, pp. 574-578.
- Gray et al., "Leases: An Efficient Fault-Tolerant Mechanism for Distributed File Cache Consistency," Proceedings of the 12th ACM Symposium on Operating Systems Principles, pp. 202-210, 1989.
- Guyennet et al., "A New Consistency Protocol Implemented in the CaliF System," IEEE, 1094-7256/97, pp. 82-87, 1997.
- Guyennet et al., "Distributed Shared Memory Layer for Cooperative Work Applications," IEEE, 0742-1303/97, pp. 72-78, 1997.
- Hoshi et al., "Allocation of the Cross-Connect Function in Leased Circuit Networks," 1992, ICC'92, conference record, SUPERCOMM/ICC '92, D a New World of Communications, IEEE International Conference, pp. 1408-1412.
- IBM Technical Disclosure Bulletin, "Local Network Monitoring to Populate Access Agent Directory," vol. 36, No. 09A, pp. 403-405, Sep. 1993.
- McGrath, "Discovery and Its Discontents: Discovery Protocols for Ubiquitous Computing," Presented at Center for Excellence in Space Data and Information Science, NASA Goddard Space Flight Center, Apr. 5, 2000.
- MUX-Elektronik, Java 1.1 Interactive Course, www.ils.se/~mux/javaic.html, 1995.
- Stevenson, "Token-Based Consistency of Replicated Servers," IEEE, CH2686-4/89/0000/0179, pp. 179-183, 1989.
- Yin et al., "Using Leases to Support Server Driven Consistency in Large-Scale Systems," Computer Services Department, University of Texas at Austin, pp. 285-294, May 26-28, 1998.
- Yin et al., "Volume Leases for Consistency in Large-Scale Systems," IEEE Transactions on Knowledge & Data Engineering, vol. 11, No. 4, pp. 563-576, Jul./Aug. 1999.
- Mullender, *Distributed Systems*, Second Edition, Addison-Wesley, 1993.
- Howard et al., Scale and Performance in a Distributed File System, ACM Transactions on Computer Systems, vol. 6, No. 1, Feb. 1988, pp. 51-81.
- Cardelli, *Obliq, A lightweight language for network objects*, Nov. 5, 1993, pp. 1-37.
- Dijkstra, Self-stabilizing Systems in Spite of Distributed Control, Communications of the ACM, vol. 17, No. 11, Nov. 1974, pp. 643-644.
- Ousterhout et al., The Sprite Network Operating System, Computer, IEEE, Feb. 1988, pp. 23-36.
- Dourish, A Divergence-Based Model of Synchrony and Distribution in Collaborative Systems, Xerox Technical Report EPC-1194-102, 1994, pp. 1-10.
- Sharrott et al., ObjectMap: Integrating High Performance Resources into a Distributed Object-oriented Environment, ICODP, 1995.
- Birrell et al., Grapevine: An Exercise in Distributed Computing, Communications of the ACM, vol. 25, No. 4, Apr. 1982, pp. 260-274.
- Transparent Network Computing*, Locus Computing Corporation, Jan. 5, 1995.
- Gray et al., Leases: An Efficient Fault-Tolerant Mechanism for Distributed File Cache Consistency, ACM, 1989, pp. 202-210.
- Lamport et al., The Byzantine Generals Problem, ACM Transactions on Programming Languages and Systems, vol. 4, No. 3, Jul. 1982, pp. 382-401.
- Dolev et al., On the Minimal Synchronism Needed for Distributed Consensus, Journal of the ACM, vol. 34, No. 1, Jan. 1987, pp. 77-97.
- Mummert et al., *Long Term Distributed File Reference Tracing: Implementation and Experience*, Carnegie Mellon University School of Computer Science, Nov. 1994, pp. 1-28.
- Gelernter et al., *Parallel Programming in Linda*, Yale University, Jan. 1985, pp. 1-21.
- Cannon et al., Adding Fault-Tolerant Transaction Processing to LINDA, Software-Practice and Experience, vol. 24(5), May 1994, pp. 449-466.
- Kambhathla et al., *Recovery with Limited Replay: Fault-Tolerant Processes in Linda*, Oregon Graduate Institute, Technical Report CSIE 90-019, Sep. 1990, pp. 1-16.
- Anderson et al., Persistent Linda: Linda + Transactions + Query Processing, Proceedings of the 13th Symposium on Fault Tolerant Systems, 1994, pp. 93-109.
- Gelernter, Generative Communication in Linda, ACM Transactions on Programming Languages and Systems, vol. 7, No. 1, Jan. 1985, pp. 80-112.
- Carriero et al., Distributed Data Structures in Linda, Principals of Programming Language, 1986, pp. 1-16.
- Pinakis, *Using Linda as the Basis of an Operating System Microkernel*, University of Western Australia, Department of Computer Science, Aug. 1993, pp. 1-165.
- Linda Database Search, Jul. 20, 1995, pp. 1-68.
- Carriero et al., *Distributed Data Structures in Linda*, Yale Research Report YALEU/DCS/RR-438, Nov. 1985.
- Agha et al., *Actorspaces: An Open Distributed Programming Paradigm*, University of Illinois, Report No. UIUCDCS-R-92-1766, Open Systems Laboratory TR No. 8, Nov. 1992, pp. 1-12.
- Ahmed et al., *A Program Building Tool for Parallel Applications*, Yale University, Dec. 1, 1993, pp. 1-23.
- Liskov et al., *Distributed Object Management in Thor*, International Workshop on Distributed Object Management, 1992, pp. 12.
- Coulouris et al., *Distributed Systems Concepts and Designs*, Second Edition, Addison-Wesley, 1994.

- Birrell et al., *Network Objects*, DEC SRC Research Report 115, Feb. 28, 1994.
- Birrell et al., *Distributed Garbage Collection for Network Objects*, DEC SRC Research Report 116, Dec. 15, 1993.
- Jaworski, *Java 1.1 Developer's Guide*, Sams.net, 1997.
- Wollrath et al., A Distributed Object Model for the Java™ System, USENIX Association, Conference on Object-Oriented Technologies and Systems, Jun. 17–21, 1996.
- Harris et al., *Proposal for a General Java Proxy Class for Distributed Systems and Other Uses*, Netscape Communications Corp., Jun. 25, 1997.
- Hamilton, Java and the Shift to Net-Centric Computing Computer, Aug. 1996, pp. 31–39.
- Chung et al., A 'Tiny' Pascal Compiler: Part 1: The P-Code Interpreter, BYTE Publications, Inc., Sep. 1978.
- Chung et al., A 'Tiny' Pascal Compiler: Part 2: The P-Compiler, BYTE Publications, Inc., Oct. 1978.
- Thompson, Regular Expression Search Algorithm, Communications of the ACM, vol. II, No. 6, p. 149 et seq., Jun. 1968.
- Mitchell et al., *Mesa Language Manual*, Xerox Corporation.
- McDaniel, *An Analysis of a Mesa Instruction Set*, Xerox Corporation, May 1982.
- Pier, *A Retrospective on the Dorado, A High-Performance Personal Computer*, Xerox Corporation, Aug. 1983.
- Pier, A Retrospective on the Dorado, A High-Performance Personal Computer, IEEE Conference Proceedings, The 10th Annual International Symposium on Computer Architecture, 1983.
- Krasner, The Smalltalk-80 Virtual Machine, BYTE Publications Inc., Aug. 1991, pp. 300–320.
- Operating Systems Review, ACM Press, vol. 27, No. 5, Dec. 1993, pp. 217–230.
- Remote Method Invocation Specification, Sun Microsystems, Inc., (1997), <http://java.sun.com/products/jdk/1.1/docs/guide/rmi/spec/rmiTOC.doc.html>.
- Alexander, et al., "Active Bridging", Proceedings of the ACM/SIGCOMM'97 Conference, Cannes, France, 9/97.
- Anonymous: "Change-Notification Service for Share Files", IBM Technical Disclosure Bulletin, vol. 36, No. 8, pp. 77–82, Aug. 1993, XP002108713, New York, US.
- Anonymous: "Resource Preemption for Priority Scheduling," Nov. 1973. IBM Technical Disclosure Bulletin, vol. 16, No. 6, p. 1931 XP002109435 New York, US.
- Beech et al., "Object Databases as Generalizations of Relational Databases," Computer Standards & Interfaces, vol. 13, Nos. 1/3, pp. 221–230, (Jan. 1991) Amsterdam, NL.
- Bertino et al., "Object-Oriented Database Management Systems: Concepts and Issues," Computer, vol. 24, No. 4, pp. 33–47, (Apr. 1991), Los Alamitos, CA.
- Betz, Mark; "Interoperable objects: laying the foundation for distributed object computing"; Dr. Dobb's Journal, vol. 19, No. 11, p. 18(13); (Oct. 1994).
- Bevan, D.I., "An Efficient Reference Counting Solution To The Distributed Garbage Collection Problem", Parallel Computing, NL, Elsevier Publishers, Amsterdam, vol. 9, No. 2, pp. 179–192, Jan. 1989.
- Birrell et al., "Implementing Remote Procedure Calls", ACM Transactions on Computer Systems, vol. 2, No. 1, Feb. 1984, pp. 39–59.
- Dave A et al: "Proxies, Application Interface, and Distributed Systems", Proceedings International Workshop on Object Orientation in Operating Systems, Sep. 24, 1992, pp. 212–220.
- Deux O et al: "The O2 System" Communications of the Association for Computing Machinery, vol. 34, No. 10, Oct. 1, 1991, pp. 34–48.
- Drexler, K. Eric, et al., "Incentive Engineering for Computational Resource Management," The Ecology of Computation, Elsevier Science Publishers B.V., 1988, pp. 231–266.
- Droms, R., "RFC 1541 Dynamic Host Configuration Protocol", [HTTP://WWW.CIS.OHIO-STATE.EDU/HTBIN/RFC/RFC1541.HTML](http://www.cis.ohio-state.edu/htbin/RFC/RFC1541.HTML), Oct. 1993, pp. 1–33.
- Emms J: "A Definition of an Access Control Systems Language" Computer Standards and Interfaces, vol. 6, No. 4, Jan. 1, 1987, pp. 443–454.
- Gosling et al., "The Java (TM) Language Specification", Addison-Wesley, 1996.
- Gottlob et al., "Extending Object-Oriented Systems with Roles," ACM Transactions on information systems, vol. 14, No. 3, pp. 268–296 (Jul. 1996).
- Guth, Rob: "JavaOne: Sun to Expand Java Distributed Computing Effort", "[HTTP://WWW.SUNWORLD.COM/SWOL-02-1998/SWOL-02-SUNSPOTS.HTML](http://WWW.SUNWORLD.COM/SWOL-02-1998/SWOL-02-SUNSPOTS.HTML)," XP-002109935, p. 1, Feb. 20, 1998.
- Hamilton et al., "Subcontract: a flexible base for distributed programming"; Proceedings of 14th Symposium of Operating System Principles; (Dec. 1993).
- Hartman, J., Manber, U., et al., Liquid Software: A new paradigm for networked systems, Technical Report 96–11, Department of Comp. Sci., Univ. of Arizona, 6/96.
- Hunt, N., "IDF: A Graphical Data Flow Programming Language for Image Processing and Computer Vision", Proceedings of the International Conference on Systems, Man, and Cybernetics, Los Angeles, Nov. 4–7, pp. 351–360, (1990).
- IBM (TM) Technical Disclosure Bulletin, "Object Location Algorithm," vol. 36, No. 09B, pp. 257–258, Sep. 1993.
- IBM (TM) Technical Disclosure Bulletin, "Retrieval of Qualified Variables Using Extendible Hashing," vol. 36, No. 12, pp. 301–303, Dec. 1993.
- Jones, Richard, et al., "Garbage Collection: Algorithms for Automatic Dynamic Memory Management," pp. 165–175, John Wiley & Sons, 1996.
- Kay, Michael H. et al., "An Overview of the Raleigh Object-Oriented Database System", ICL Technical Journal, vol. 7, No. 4, pp. 780–798, (Nov. 1991), Oxford, GB.
- Kougiouris et al.; "Support for Space Efficient Object Invocation in Spring"; (Sep. 1994).
- Lindholm et al., "The Java (TM) Virtual Machine Specification", Addison Wesley, 1996.
- Mitchell et al.; "An Overview of the Spring System"; (Feb. 1994).
- Riggs Roger et al., "Picking State in the Java (TM) System," USENIX Association Conference on Object-Oriented Technologies and Systems, XP-002112719, Jun. 17–21, 1996, pp. 241–250.
- Rosenberry et al., "Understanding DCE"; Chapters 1–3, 6; (1992).
- Waldo J et al: "Events in an RPC based distributed system" Proceedings of the 1995 USENIX Technical Conference, Proceedings USENIX Winter 1995 Technical Conference, New Orleans, LA. USA, Jan. 16–20, 1995, pp. 131–142.
- Wilson, P.R., et al., "Design of the Opportunistic Garbage Collector," Proceedings of the Object Oriented Programming Systems Languages And Applications Conference, New Orleans, vol. 24, No. 10, Oct. 1989.

Wu, Xuequn, "A Type system for an Object-Oriented Database System," Proceedings of the International Computer Software and Applications Conference (COMPSAC), pp. 333-338, Sep. 11-13, 1991, Tokyo, Japan.

Yemini, Y. and S. da Silva, "Towards Programmable Networks" IFIP/IEEE International Workshop on Distributed Systems: Operations and Management, L'Aquila, Italy, 10/96.

Birrell et al., Implementing Remote Procedure Calls, ACM Transactions on Computer Systems, vol. 2, No. 1, Feb. 1984, pp. 39-59.

IBM: Somobjects Developer Toolkit Users Guide, Version 2.1, "Chapter 6 Distributed SOM (DSOM)," pp 6-1-6-90, Oct. 1994.

Orfali R. et al., "The Essential Distributed Objects Survival Guide," Chapter 11: Corba Commercial ORBs, John Wiley & Sons, Inc., (1996).

Venners, B., "Jini Technology, Out of the Box", Javaworld, Online!, pp 1-4, Dec. 1998.

* cited by examiner




FIG. 1




FIG. 2





FIG. 3

**FIG. 4**



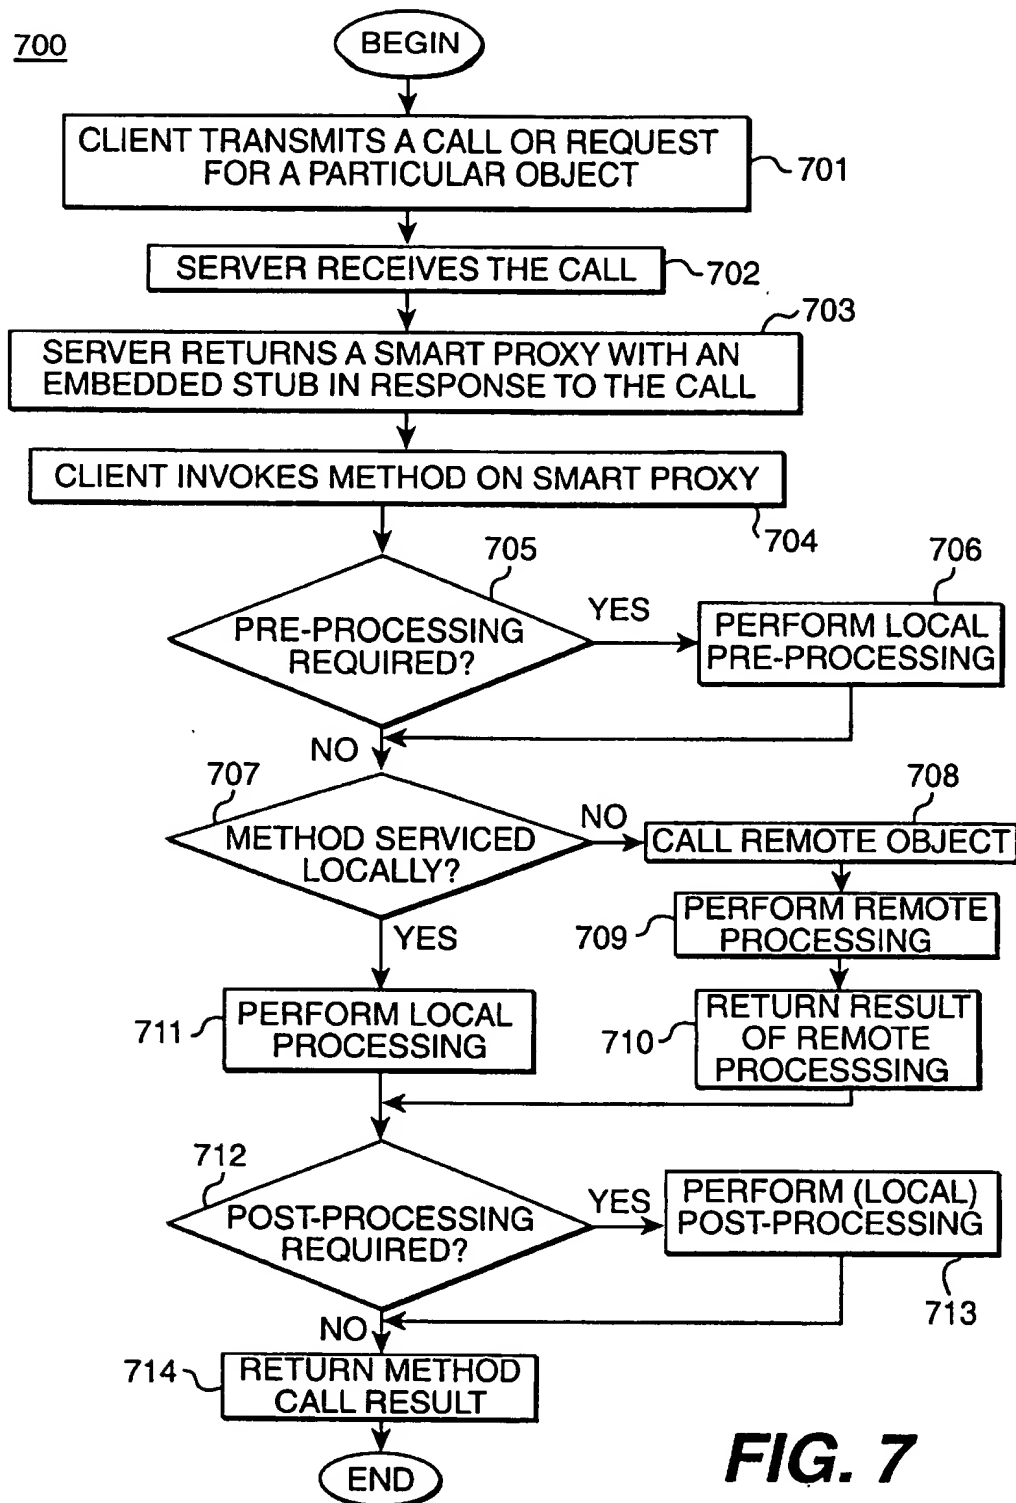
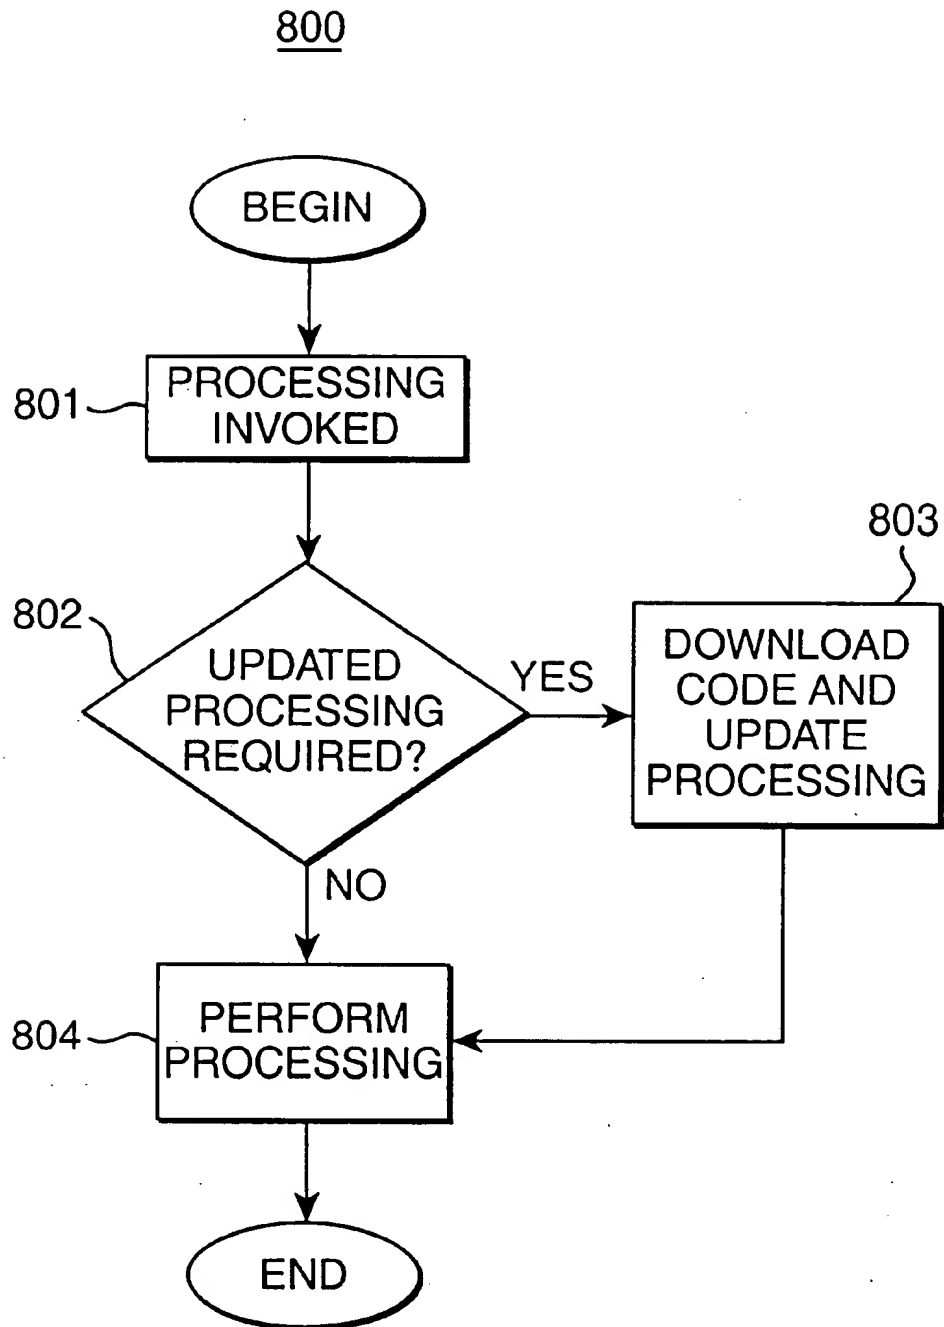


FIG. 5

**FIG. 6**

**FIG. 7**

**FIG. 8**

1

DOWNLOADABLE SMART PROXIES FOR PERFORMING PROCESSING ASSOCIATED WITH A REMOTE PROCEDURE CALL IN A DISTRIBUTED SYSTEM

REFERENCE TO RELATED APPLICATIONS

Provisional U.S. patent application No. 60/076,048, entitled "Distributed Computing System," filed on Feb. 26, 1998.

U.S. patent application Ser. No. 09/044,923, entitled "Method and System for Leasing Storage," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,838, entitled "Method, Apparatus, and Product for Leasing of Delegation Certificates in a Distributed System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,834, entitled "Method, Apparatus and Product for Leasing of Group Membership in a Distributed System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,916, entitled "Leasing for Failure Detection," filed on the same date herewith.

U.S. patent application Ser. No. 09/144,933, entitled "Method for Transporting Behavior in Event Based System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,919, entitled "Deferred Reconstruction of Objects and Remote Loading for Event Notification in a Distributed System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,938, entitled "Methods and Apparatus for Remote Method Invocation," filed on the same date herewith.

U.S. patent application Ser. No. 09/045,652, entitled "Method and System for Deterministic Hashes to Identify Remote Methods," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,790, entitled "Method and Apparatus for Determining Status of Remote Objects in a Distributed System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,917, entitled "Suspension and Continuation of Remote Methods," and filed on the same date herewith.

U.S. patent application Ser. No. 09/044,835, entitled "Method and System for Multi-Entry and Multi-Template Matching in a Database," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,839, entitled "Method and System for In-Place Modifications in a Database," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,945, entitled "Method and System for Typesafe Attribute Matching in a Database," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,931, entitled "Dynamic Lookup Service in a Distributed System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,939, entitled "Apparatus and Method for Providing Downloadable Code for Use in Communicating with a Device in a Distributed System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,826, entitled "Method and System for Facilitating Access to a Lookup Service," filed on the same date herewith.

U.S. patent application No. 09/030,840, entitled "Method and Apparatus for Dynamic Distributed Computing Over a Network," and filed on Feb. 26, 1998.

2

U.S. patent application Ser. No. 09/044,932, entitled "Apparatus and Method for Dynamically Verifying Information in a Distributed System," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,936, entitled "An Interactive Design Tool for Persistent Shared Memory Spaces," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,934, entitled "Polymorphic TokenBased Control," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,915, entitled "Stack-Based Access Control," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,944, entitled "Stack-Based Security Requirements," filed on the same date herewith.

U.S. patent application Ser. No. 09/044,837, entitled "Per-Method Designation of Security Requirements," filed on the same date herewith.

FIELD OF THE INVENTION

The present invention relates to a system and method for transmitting objects between machines in a distributed system and more particularly relates to transmission of a representation of a remote object including code for local processing.

BACKGROUND OF THE INVENTION

Distributed programs which concentrate on point-to-point data transmission can often be adequately and efficiently handled using special-purpose protocols for remote terminal access and file transfer. Such protocols are tailored specifically to the one program and do not provide a foundation on which to build a variety of distributed programs (e.g., distributed operating systems, electronic mail systems, computer conferencing systems, etc.).

While conventional transport services can be used as the basis for building distributed programs, these services exhibit many organizational problems, such as the use of different data types in different machines, lack of facilities for synchronization, and no provision for a simple programming paradigm.

Distributed systems usually contain a number of different types of machines interconnected by communications networks. Each machine has its own internal data types, its own address alignment rules, and its own operating system. This heterogeneity causes problems when building distributed systems. As a result, program developers must include in programs developed for such heterogeneous distributed systems the capability of ensuring that information is handled and interpreted consistently in different machines.

However, one simplification is afforded by noting that a large proportion of programs use a request and response interaction between processes where the initiator (i.e., program initiating a communication) is blocked out until the response is returned and is thus idle during this time. This can be modeled by a procedure call mechanism between processes. One such mechanism is referred to as the remote procedure call (RPC).

RPC is a mechanism for providing synchronized communication between two processes (e.g., program, applet, etc.) running on the same machine or different machines. In a simple case, one process, e.g., a client program, sends a message to another process, e.g., a server program. In this case, it is not necessary for the processes to be synchronized

either when the a message is sent or received. It is possible for the client program to transmit the message and then begin a new activity, or for the server program's environment to buffer the incoming message until the server program is ready to process a new message.

RPC, however, imposes constraints on synchronism because it closely models the local procedure call, which requires passing parameters in one direction, blocking the calling process (i.e., the client program) until the called procedure of the server program is complete, and then returning a response. RPC thus involves two message transfers, and the synchronization of the two processes for the duration of the call.

The R-PC mechanism is usually implemented in two processing parts using the local procedure call paradigm, one part being on the client side and the other part being on the server side. Both of these parts will be described below with reference to FIG. 1.

FIG. 1 is a diagram illustrating the flow of call information using an RPC mechanism. As shown in FIG. 1, a client program 100 issues a call (step 102). The RPC mechanism 101 then packs the call as arguments of a call packet (step 103), which the RPC mechanism 101 then transmits to a server program 109 (step 104). The call packet also contains information to identify the client program 100 that first sent the call. After the call packet is transmitted (step 104), the RPC mechanism 101 enters a wait state during which it waits for a response from the server program 109.

The RPC mechanism 108 for the server program 109 (which may be the same RPC mechanism as the RPC mechanism 101 when the server program 109 is on the same platform as the client program 100) receives the call packet (step 110), unpacks the arguments of the call from the call packet (step 111), identifies, using the call information, the server program 109 to which the call was addressed, and provides the call arguments to the server program 109.

The server program receives the call (step 112), processes the call by invoking the appropriate procedure (step 115), and returns a response to the RPC mechanism 108 (step 116). The RPC mechanism 108 then packs the response in a response packet (step 114) and transmits it to the client program 100 (step 113).

Receiving the response packet (step 107) triggers the RPC mechanism 101 to exit the wait state and unpack the response from the response packet (step 106). RPC 101 then provides the response to the client program 100 in response to the call (step 105). This is the process flow of the typical RPC mechanism modeled after the local procedure call paradigm. Since the RPC mechanism uses the local procedure call paradigm, the client program 100 is blocked at the call until a response is received. Thus, the client program 100 does not continue with its own processing after sending the call; rather, it waits for a response from the server program 109.

The Java™ programming language is an object-oriented programming language that is typically compiled into a platform-independent format, using a bytecode instruction set, which can be executed on any platform supporting the Java virtual machine (JVM). This language is described, for example, in a text entitled "The Java Language Specification" by James Gosling, Bill Joy, and Guy Steele, Addison-Wesley, 1996, which is incorporated herein by reference. The JVM is described, for example, in a text entitled "The Java Virtual Machine Specification," by Tim Lindholm and Frank Yellin, Addison Wesley, 1996, which is incorporated herein by reference. Java and Java-based trademarks are

trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Because the JVM may be implemented on any type of platform, implementing distributed programs using the JVM significantly reduces the difficulties associated with developing programs for heterogeneous distributed systems. Moreover, the JVM uses a Java remote method invocation system (RMI) that enables communication among programs of the system. RMI is explained in, for example, the following document, which is incorporated herein by reference: Remote Method Invocation Specification, Sun Microsystems, Inc. (1997), which is available via universal resource locator (URL)

<http://www.javasoft.com/products/jdk/1.1/docs/guide/rmi/spec/rmiTOC.doc.html>.

FIG. 2 is a diagram illustrating the flow of objects in an object-oriented distributed system 200 including machines 201 and 202 for transmitting and receiving method invocations using the JVM. In system 200, machine 201 uses RMI 205 for responding to a call for object 203 by converting the object into a byte stream 207, including an identification of the type of object transmitted and data constituting the object. While machine 201 is responding to the call for object 203, a process running on the same or another machine in system 200 may continue operation without waiting for a response to its request.

Machine 202 receives the byte stream 207. Using RMI 206, machine 202 automatically converts it into the corresponding object 204, which is a copy of object 203 and which makes the object available for use by a program executing on machine 202. Machine 202 may also transmit the object to another machine by first converting the object into a byte stream and then sending it to the third machine, which also automatically converts the byte stream into the corresponding object.

The communication between these machines sometimes involves, for example, repeated calls for the same information. These calls are made to a local proxy, which acts as a surrogate for the remote object in the address space of the client. Such a proxy will service the call by making a network request to the server object. Repeated calls to the same server object through a proxy can generate considerable network traffic, increasing the time and expense of obtaining the information. Accordingly, a need exists for a technique that reduces the amount of network communication in, for example, such a case.

SUMMARY OF THE INVENTION

A method consistent with the present invention transmits a request for a particular object. A response to the request is received, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the object, local to the requesting object, using the representation.

Another method consistent with the present invention receives at a machine a request for a particular object. A response to the request is transmitted, the response including first code for constructing a representation of the object and including an indication of second code for processing, such that the construction creates an object for processing calls to the object, local to the requesting object, using the representation.

An apparatus consistent with the present invention transmits a request for a particular object. The apparatus receives a response to the request, the response including code used to construct a representation of the requested object, the

5

construction creating an object for processing calls to the object, local to the requesting object, using the representation.

Another apparatus consistent with the present invention receives at a machine a request for a particular object. The apparatus transmits a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing, such that the construction creates an object for processing calls to the object, local to the requesting object, using the representation.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings are incorporated in and constitute a part of this specification and, together with the description, explain the advantages and principles of the invention. In the drawings,

FIG. 1 is a diagram illustrating the flow of call information using an RPC mechanism;

FIG. 2 is a diagram illustrating the transmission of objects in an object-oriented distributed system;

FIG. 3 is a diagram of an exemplary distributed processing system that can be used in an implementation consistent with the present invention;

FIG. 4 is a diagram of an exemplary distributed system infrastructure;

FIG. 5 is a diagram of a computer in a distributed system infrastructure shown in FIG. 4;

FIG. 6 is a block diagram of a distributed network for use in downloading smart proxies;

FIG. 7 is a flow chart of a process for downloading smart proxies within, for example, the distributed network shown in FIG. 6; and

FIG. 8 is a flow chart of a process for changing the processing performed by a smart proxy.

DETAILED DESCRIPTION

Overview

Instead of receiving a proxy that only makes network requests to the object for which it is a surrogate, a machine in a distributed system receives a smart proxy. Such a proxy can respond to calls on the object for which it is a surrogate without making any network calls to increase program efficiency, or perform processing before making a network call or after the completion of the network call to increase program functionality. The term proxy generally refers to code or other mechanism used to act as a surrogate for a remote object in the address space of a machine.

Systems transferring stubs and associated smart proxies may use a variant of an RPC or RMI, passing arguments and return values from one process to another process each of which may be on different machines. The term "machine" is used in this context to refer to a physical machine or a virtual machine. Multiple virtual machines may exist on the same physical machine. Examples of RPC systems include distributed computed environment (DCE) RPC and Microsoft distributed common object model (DCOM) RPC. A memory stores the stub and associated smart proxy, and this memory may include secondary sources such as a disk or receiving objects from the Internet.

Distributed Processing System

FIG. 3 illustrates an exemplary distributed processing system 300 which can be used in an implementation con-

6

sistent with the present invention. In FIG. 3, distributed processing system 300 contains three independent and heterogeneous platforms 301, 302, and 303 connected in a network configuration represented by network cloud 319.

The composition and protocol of the network configuration represented by cloud 319 is not important as long as it allows for communication of the information between platforms 301, 302 and 303. In addition, the use of just three platforms is merely for illustration and does not limit an implementation consistent with the present invention to the use of a particular number of platforms. Further, the specific network architecture is not crucial to embodiments consistent with this invention. For example, another network architecture that could be used in an implementation consistent with this invention would employ one platform as a network controller to which all the other platforms would be connected.

In the implementation of distributed processing system 300, platforms 301, 302 and 303 each include a processor 316, 317, and 318 respectively, and a memory, 304, 305, and 306, respectively. Included within each memory 304, 305, and 306, are applications 307, 308, and 309, respectively, operating systems 310, 311, and 312, respectively, and RMI components 313, 314, and 315, respectively.

Applications 307, 308, and 309 can be applications or programs that are either previously written and modified to work with, or that are specially written to take advantage of, the services offered by an implementation consistent with the present invention. Applications 307, 308, and 309 invoke operations to be performed in accordance with an implementation consistent with this invention.

Operating systems 310, 311, and 312 are typically standard operating systems tied to the corresponding processors 316, 317, and 318, respectively. The platforms 301, 302, and 303 can be heterogeneous. For example, platform 301 has an UltraSparc® microprocessor manufactured by Sun Microsystems, Inc. as processor 316 and uses a Solaris® operating system 310. Platform 302 has a MIPS microprocessor manufactured by Silicon Graphics Corp. as processor 317 and uses a Unix operating system 311. Finally, platform 303 has a Pentium microprocessor manufactured by Intel Corp. as processor 318 and uses a Microsoft Windows 95 operating system 312. An implementation consistent with the present invention is not so limited and could accommodate homogenous platforms as well.

Sun, Sun Microsystems, Solaris, Java, and the Sun Logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UltraSparc and all other SPARC trademarks are used under license and are trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

Memories 304, 305, and 306 serve several functions, such as general storage for the associated platform. Another function is to store applications 307, 308, and 309, RMI components 313, 314, and 315, and operating systems 310, 311, and 312 during execution by the respective processor 316, 317, and 318. In addition, portions of memories 304, 305, and 306 may constitute shared memory available to all of the platforms 301, 302, and 303 in network 319. Note that RMI components 313, 314, and 315 operate in conjunction with a JVM, which is not shown for the purpose of simplifying the figure.

Distributed System Infrastructure

Systems and methods consistent with the present invention may also operate within a particular distributed system

400, which will be described with reference to FIGS. 4 and 5. This distributed system 400 is comprised of various components, including hardware and software, to (1) allow users of the system to share services and resources over a network of many devices; (2) provide programmers with tools and programming patterns that allow development of robust, secured distributed systems; and (3) simplify the task of administering the distributed system. To accomplish these goals, distributed system 400 utilizes the Java programming environment to allow both code and data to be moved from device to device in a seamless manner. Accordingly, distributed system 400 is layered on top of the Java programming environment and exploits the characteristics of this environment, including the security offered by it and the strong typing provided by it.

In distributed system 400 of FIGS. 4 and 5, different computers and devices are federated into what appears to the user to be a single system. By appearing as a single system, distributed system 400 provides the simplicity of access and the power of sharing that can be provided by a single system without giving up the flexibility and personalized response of a personal computer or workstation. Distributed system 400 may contain thousands of devices operated by users who are geographically disperse, but who agree on basic notions of trust, administration, and policy.

Within an exemplary distributed system are various logical groupings of services provided by one or more devices, and each such logical grouping is known as a Djinn. A "service" refers to a resource, data, or functionality that can be accessed by a user, program, device, or another service and that can be computational, storage related, communication related, or related to providing access to another user. Examples of services provided as part of a Djinn include devices, such as printers, displays, and disks; software, such as programs or utilities; information, such as databases and files; and users of the system.

Both users and devices may join a Djinn. When joining a Djinn, the user or device adds zero or more services to the Djinn and may access, subject to security constraints, any one of the services it contains. Thus, devices and users federate into a Djinn to share access to its services. The services of the Djinn appear programmatically as objects of the Java programming environment, which may include other objects, software components written in different programming languages, or hardware devices. A service has an interface defining the operations that can be requested of that service, and the type of the service determines the interfaces that make up that service.

Distributed system 400 is comprised of computer 402, a computer 404, and a device 406 interconnected by a network 408. Device 406 may be any of a number of devices, such as a printer, fax machine, storage device, computer, or other devices. Network 408 may be a local area network, wide area network, or the Internet. Although only two computers and one device are depicted as comprising distributed system 400, one skilled in the art will appreciate that distributed system 400 may include additional computers or devices.

FIG. 5 depicts computer 402 in greater detail to show a number of the software components of distributed system 400. One skilled in the art will appreciate that computer 404 or device 406 may be similarly configured. Computer 402 includes a memory 502, a secondary storage device 504, a central processing unit (CPU) 506, an input device 508, and a video display 510. Memory 502 includes a lookup service 512, a discovery server 514, and a Java runtime system 516. The Java runtime system 516 includes the Java RMI system

518 and a JVM 520. Secondary storage device 504 includes a Java space 522.

As mentioned above, distributed system 400 is based on the Java programming environment and thus makes use of the Java runtime system 516. The Java runtime system 516 includes the Java API libraries, allowing programs running on top of the Java runtime system to access, in a platform-independent manner, various system functions, including windowing capabilities and networking capabilities of the host operating system. Since the Java API libraries provide a single common API across all operating systems to which the Java runtime system is ported, the programs running on top of a Java runtime system run in a platform-independent manner, regardless of the operating system or hardware configuration of the host platform. The Java runtime system 516 is provided as part of the Java software development kit available from Sun Microsystems, Inc. of Mountain View, Calif.

JVM 520 also facilitates platform independence. JVM 520 acts like an abstract computing machine, receiving instructions from programs in the form of bytecodes and interpreting these bytecodes by dynamically converting them into a form for execution, such as object code, and executing them. RMI 518 facilitates remote method invocation by allowing objects executing on one computer or device to invoke methods of an object on another computer or device. Both RMI and the JVM are also provided as part of the Java software development

Lookup service 512 defines the services that are available for a particular Djinn. That is, there may be more than one Djinn and, consequently, more than one lookup service within distributed system 400. Lookup service 512 contains one object for each service within the Djinn, and each object contains various methods that facilitate access to the corresponding service. Lookup service 512 is described in U.S. patent application entitled "Method and System for Facilitating Access to a Lookup Service," which was previously incorporated herein by reference.

Discovery server 514 detects when a new device is added to distributed system 400, during a process known as boot and join (or discovery), and when such a new device is detected, the discovery server passes a reference to lookup service 512 to the new device so that the new device may register its services with the lookup service and become a member of the Djinn. After registration, the new device becomes a member of the Djinn, and as a result, it may access all the services contained in lookup service 512. The process of boot and join is described in U.S. patent application entitled "Apparatus and Method for providing Downloadable Code for Use in Communicating with a Device in a Distributed System," which was previously incorporated herein by reference.

A Java space 522 is an object repository used by programs within distributed system 400 to store objects. Programs use a Java space 522 to store objects persistently as well as to make them accessible to other devices within distributed system 400. Java spaces are described in U.S. patent application Ser. No. 08/971,529, entitled "Database System Employing Polymorphic Entry and Entry Matching," assigned to a common assignee, and filed on Nov. 17, 1997, which is incorporated herein by reference. One skilled in the art will appreciate that an exemplary distributed system 400 may contain many lookup services, discovery servers, and Java spaces.

Data Flow in a Distributed Processing System

FIG. 6 is a block diagram of an object-oriented distributed network 600 connecting machines 601 and 606, such as

computers or virtual machines executing on one or more computers, or the machines described with reference to FIGS. 3, 4, and 5. Network 600 transmits proxies, some of which may be smart proxies. A smart proxy includes code for performing processing associated with a call. For example, a smart proxy may perform a caching operation for read-only data for later reference. When a call is made for that data, the smart proxy may obtain it locally and provide it to a user without making another call for the data, which may occur transparent to the user. An example of such read-only data is a particular installation time. The first time a call is made for the installation time, for example, a smart proxy locally caches that value, and when a subsequent call is made for the installation time, the smart proxy locally retrieves the value.

Another example of smart proxy processing involves use of a serialized object for transmitting data to a data bank storing information. In this example, a call is made to a smart proxy, which receives an object, serializes the object on the client machine into an array of bytes, and transmits the array of bytes to a server. The server only stores the serialized object, avoiding the requirement to download code, and it provides a key for the object to the client machine. When the client machine wants to retrieve the data, the smart proxy transmits the key to the server, receives in response the serialized object, reconstructs the object, and provides it to the user.

Other examples of uses of smart proxies include processing for debugging, call logging, and monitoring system performance. Another example involves the use of a smart proxy for local data verification, as explained in the U.S. patent application filed on the same date herewith, assigned to a common assignee, and entitled "Apparatus and Method for Dynamically Verifying Information in a Distributed System," which is incorporated herein by reference. Many other uses for smart proxies are possible for performing processing associated with a call.

Network 600 includes a client machine 601 containing RMI 602 and associated code 603. A server machine 606 includes RMI 607 and remote object 608. In operation, RMI 602 transmits a call or request 609 to RMI 607, requesting a particular stub object. RMI 607 returns a response 610 including requested stub 605 embedded within a smart proxy 604. The response may be transmitted as a stream. Streams used in the Java programming language, including input and output streams, are known in the art and an explanation, which is incorporated herein by reference, appears in, for example, a text entitled "The Java Tutorial: Object-Oriented

Programming for the Internet," pp. 325-53, by Mary Campione and Kathy Walrath, Addison-Wesley, 1996.

The response may include information so that client machine 601 can reconstruct the stub object in smart proxy 604. When a set of object types is limited and is the same on machines 601 and 606, a receiving machine typically requires the object's state and a description of its type because the object's code is already present on all network machines. Alternatively, machine 606 uses RMI 607 to provide more flexibility, allowing code to be moved when necessary along with information or the object's state and type. Additionally, a transmitting machine may include in the object an identification of the type of object transmitted, the data constituting the state of the object, and a network-accessible location in the form of a URL for code that is associated with the object. URLs are known in the art and an explanation, which is incorporated herein by reference,

appears in, for example, a text entitled "The Java Tutorial: Object-Oriented Programming for the Internet," pp. 494-507, by Mary Campione and Kathy Walrath, Addison-Wesley, 1996.

When client machine 601 receives response 610, it identifies the type of transmitted object. Machine 601 contains its own RMI 602 and code 603 for processing of objects, and it may create stub object 605 using the object type, the state information, and code for the object. If code for the object is not resident or available on machine 601 and the stub object does not contain the code, RMI 602 may use a URL from the object to locate the code and transfer a copy of the code to client machine 601. Because the code is bytecodes and is therefore portable, client machine 601 can load the code into RMI 602 to reconstruct the object. Thus, client machine 601 can reconstruct an object of the appropriate type even if that kind of object has not been present on the machine before.

When creating stub object 605, RMI 602 does not necessarily know that the stub is itself a smart proxy 604. Smart proxy 604 may perform processing at client machine 601 before or after response 610 and may supply all processing without resorting to call 609 to the object for which the proxy acts. Therefore, smart proxy 604 may perform all processing locally when client machine 601 makes a call or request 611 to invoke a method on smart proxy 604. These proxies are downloadable by the same methods as disclosed in U.S. patent application Ser. No. 08/950,756, filed on Oct. 15, 1997, and entitled "Deferred Reconstruction of Objects and Remote Loading in a Distributed System," which is incorporated herein by reference.

Transmission of Smart Proxies

FIG. 7 is a flow chart of a process 700 for downloading and using smart proxies within, for example, the distributed network shown in FIG. 6. A client machine transmits a call or request for a particular object (step 701), and a server machine receives the call (step 702). In response, the server machine returns a smart proxy with an embedded stub (step 703), and the proxy acts as a representation of the requested object. After receiving the smart proxy, the client machine invokes a method on it (step 704). According to the code within the smart proxy, the client machine containing the smart proxy determines if preprocessing is required (step 705). If so, the processing is performed locally by the client machine using the smart proxy (step 706).

The client machine then determines if the method called on the smart proxy may be serviced locally (step 707). If so, the client machine performs the local processing for the call (step 711). If not, the client machine calls the remote object (step 708). The remote processing is performed (step 709), and the result of the remote processing is returned to the client machine (step 710).

The client machine determines, according to code in the smart proxy, if post-processing as a result of the call is required (step 712). If so, it locally performs the post-processing using code in the smart proxy (step 713). The smart proxy then returns the method call result (step 714) in response to the call on the smart proxy in step 704.

FIG. 8 is a flow chart of a process 800 for changing the processing performed by a smart proxy. When processing is invoked (step 801), a client machine determines if updated processing is required (step 802). Such information may be contained within the smart proxy itself in that it may determine when or under what particular circumstances it requires updated processing code. If updated processing is

11

required, the code for that processing is downloaded and the smart proxy is updated at the client machine to perform that processing (step 803). The smart proxy then performs at the client machine the processing according to the updated code (step 804).

Machines implementing the steps shown in FIGS. 7 and 8 may include computer processors for performing the functions, as shown in FIGS. 3, 4, 5, and 6. They may include modules or programs configured to cause the processors to perform the above functions. They may also include computer program products stored in a memory. The computer program products may include a computer-readable medium or media having computer-readable code embodied therein for causing the machines to perform functions described above. The computer-readable media may include computer data signals embodied in a carrier wave and representing sequences of instructions which, when executed by a processor, cause the processor to securely address a peripheral device at an absolute address by performing the method described in this specification. The media may also include a data structure for use in performing the method described in this specification.

Although the illustrative embodiments of the systems consistent with the present invention are described with reference to a computer system implementing the Java programming language on the JVM specification, the invention is equally applicable to other computer systems processing code from different programming languages. Specifically, the invention may be implemented with both object-oriented and nonobject-oriented programming systems. In addition, although an embodiment consistent with the present invention has been described as operating in the Java programming environment, one skilled in the art will appreciate that the present invention can be used in other programming environments as well.

While the present invention has been described in connection with an exemplary embodiment, it will be understood that many modifications will be readily apparent to those skilled in the art, and this application is intended to cover any adaptations or variations thereof. For example, different labels or definitions for the smart proxies may be used without departing from the scope of the invention. This invention should be limited only by the claims and equivalents thereof.

What is claimed is:

1. A method for receiving objects in a distributed system comprised of multiple machines, comprising:
 - transmitting a request for a particular object; and
 - receiving a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to the requesting object, using the representation.
2. The method of claim 1, further including using the code to construct the representation.
3. The method of claim 2 wherein the using includes downloading code to update the processing.
4. A method for transmitting objects in a distributed system comprised of multiple machines, comprising:
 - receiving at a machine a request for a particular object; and
 - transmitting a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing such that the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.

12

5. An apparatus for receiving objects in a distributed system comprised of multiple machines, comprising:

- a module configured to transmit a request for a particular object; and
- a module configured to receive a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to the requesting object, using the representation.

6. The apparatus of claim 5, further including

- a module configured to use the code to construct the representation.

7. The apparatus of claim 6 wherein the use module includes

- a module configured to download code to update the processing.

8. An apparatus for transmitting objects in a distributed system comprised of multiple machines, comprising:

- a module configured to receive at a machine a request for a particular object; and
- a module configured to transmit a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing such that the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.

9. A system for transmitting objects in a distributed system comprised of multiple machines, comprising:

- a first machine;
- a second machine;
- a network connecting the first machine with the second machine; and

an apparatus for receiving objects, the apparatus including:

- a module configured to transmit a request for a particular object; and
- a module configured to receive a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to the requesting object, using the representation.

10. The system of claim 9, further including

- a module configured to use the code to construct the representation.

11. The system of claim 10 wherein the use module includes

- a module configured to download code to update the processing.

12. A system for transmitting objects in a distributed system comprised of multiple machines, comprising:

- a first machine;
- a second machine;
- a network connecting the first machine with the second machine; and an apparatus for transmitting objects, the apparatus including:
 - a module configured to receive at a machine a request for a particular object; and
 - a module configured to transmit a response to the request, the response including first code for constructing a representation of the object and including an indication of second code for processing such that

13

the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.

13. A computer program product, comprising:

a computer-readable medium containing instructions for controlling a computer system to perform a method, the method including:

transmitting a request for a particular object; and

receiving a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to the requesting object, using the representation.

14. The computer program product of claim 13, further including

using the code to construct the representation.

15. The computer program product of claim 14 wherein the using includes

downloading code to update the processing.

16. A computer program product, comprising:

a computer-readable medium containing instructions for controlling a computer system to perform a method, the method including:

receiving at a machine a request for a particular object; and

transmitting a response to the request, the response including first code for constructing a representation of

14

the object and including an indication of second code for processing such that the construction creates an object for processing calls to the particular object, local to the requesting object, using the representation.

17. An article of manufacture specifying a representation of an object stored in a computer-readable storage medium and capable of electronic transmission between machines in a distributed system, the article of manufacture comprising:

an object to be transmitted from a first machine to a second machine in a response to a request, the object including an indication of code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to the requesting object, using the representation.

18. An apparatus for receiving objects in a distributed system comprised of multiple machines, comprising:

means for transmitting a request for a particular object; and

means for receiving a response to the request, the response including code used to construct a representation of the requested object, the construction creating an object for processing calls to the particular object, local to the requesting object, using the representation.

* * * * *